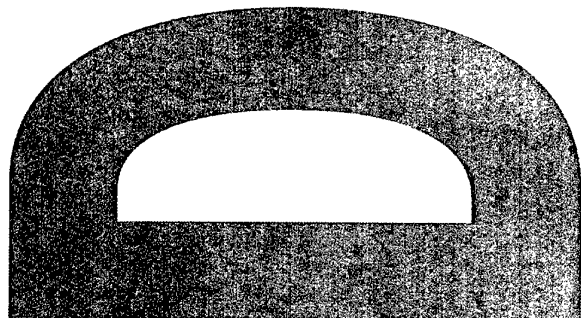
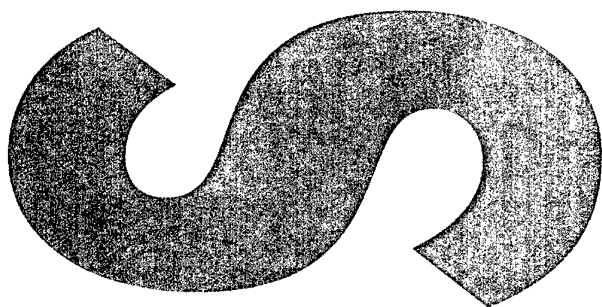
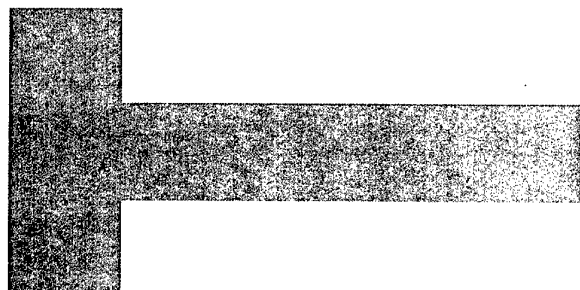
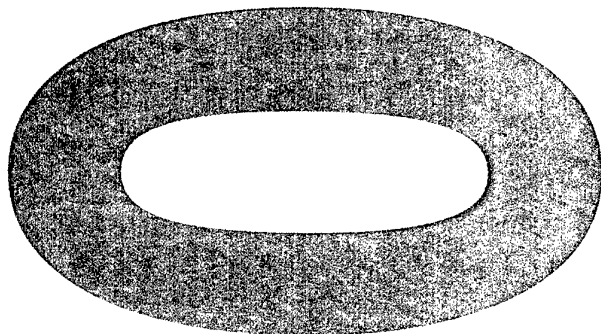




Australian Government

Department of Defence
Defence Science and
Technology Organisation



Enhancing Imagery to Improve the Visibility of Detail

Robert Whatmough

DSTO-TR-1557

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20040617 032



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Enhancing Imagery to Improve the Visibility of Detail

Robert Whatmough

Intelligence, Surveillance and Reconnaissance Division
Information Sciences Laboratory

DSTO-TR-1557

ABSTRACT

Imagery from surveillance sensors is often not suitable for immediate viewing, because the number of grey levels or colours used is much larger or smaller than the number available on the display device. This Report considers some old and new techniques for making imagery ready to view, so that both broad features and fine detail are visible. The basic techniques often give poor results when they encounter the peculiarities of particular sensors or scenes; in these cases variations are available that overcome the problems at a cost in computing time. Examples of the application of techniques and variations to airborne sensor imagery are given.

RELEASE LIMITATION

Approved for public release

AQ F04-09-0836

Published by

*DSTO Information Sciences Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2004
AR-013-039
February 2004*

APPROVED FOR PUBLIC RELEASE

Enhancing Imagery to Improve the Visibility of Detail

Executive Summary

Surveillance sensors often produce imagery that is not ready to view because the number of grey levels or colours used is much smaller or larger than the number available on computer display devices. Some adjustment is needed to ensure that the imagery makes good use of the display, with appropriate contrast over the whole image and in the finer details, so that a human analyst can make full and rapid use of the information available.

Simple scaling of the grey values of monochrome imagery is usually inadequate unless the imagery already makes good use of the range. Normalisation, which stretches the range actually used to the maximum, is better but is easily misled by sensor or processing artefacts that produce bright or dark areas with greater range than the useful data. Further adjustments to normalisation can overcome the problems, but some analyst interpretation will still be needed, at least for representative images in a series.

Histogram modification is a time-honoured technique for improving contrast when parts of the range of grey levels are unequally used. It may, however, cause over-amplification of noise or excessive loss of contrast for certain scenes. Modifications are available to control these problems.

Local contrast enhancement, in which the image contrast is modified differently in different areas, can be performed in several ways. Care is needed to avoid loss of intelligibility of the image, the amplification of noise and the introduction of distracting artefacts. Increasing the contrast in bright areas is not usually performed in this process, but it turns out to be a useful part of it. Some welcome sharpening of images results as a by-product of this enhancement.

The theory behind the above enhancement techniques often does not extend readily to colour imagery, but the results can easily be applied to it. The television engineer's distinction between luminance (brightness) and chrominance (colour) must be respected to avoid new artefacts associated with colour.

Some allowance may need to be made for the characteristics of the device used for display, especially if it is a monitor.

This report presents numerous examples of enhancement of images from airborne sensors, and the techniques described would be useful in ground station software systems for such sensors.

Authors

Robert Whatmough

Intelligence, Surveillance and Reconnaissance
Division

Robert Whatmough was employed as Cadet Defence Science at the Weapons Research Establishment, Department of Supply, later DSTO. He was awarded the degree of B Sc with honours in 1969.

Until 1985 he was Experimental Officer, Scientific Officer then Research Scientist in Computing Services Group. His work included mathematical data handling techniques, random number generation, printed circuit board design, simulation of batch computer operations, curve fitting and smoothing, time-critical computing, thermal modelling, computer graphics and display, remote sensing and image processing, and assistance of computer users with complex computing problems.

Since 1986 he has worked in various Divisions in fields related to image processing. These have included restoration, enhancement and classification of visual, multispectral and synthetic aperture radar images, geographic information systems, prediction of oblique aerial and navigation radar images, matching models to objects in images, shape inference from perspective distortion, registration and mosaicking of aerial images, and enhancement of video sequences.

Contents

1. INTRODUCTION.....	1
2. LINEAR CONTRAST ENHANCEMENT	1
2.1 Scaling	1
2.2 Normalisation	2
2.3 Trimmed normalisation	2
3. ENHANCEMENT BY HISTOGRAM MODIFICATION.....	3
3.1 General approach	3
3.2 Histogram equalisation.....	3
3.3 Bounded histogram equalisation	4
3.4 Optimal bounded histogram equalisation.....	4
4. LOCAL CONTRAST ENHANCEMENT	7
4.1 Possible approaches.....	7
4.2 The ACE method	8
5. APPLICATION TO COLOUR IMAGERY.....	10
5.1 Colour representations	10
5.2 Extension of grey-level methods	11
6. NON-LINEAR EFFECTS OF DISPLAY DEVICES.....	12
7. CONCLUSIONS	13

1. Introduction

Surveillance sensors often produce imagery that is not ready to view because the number of grey levels or colours used is much smaller or larger than the number available on computer display devices. For example, many IR sensors record far more than the 256 levels normally displayed, to allow for the great range of radiances that might be encountered; the number of grey levels used in a particular scene may be far smaller than the number allowed, even though there are many objects of interest in the scene.

An analyst who views an image from such a sensor may be interested in seeing anything present that can be distinguished from random noise and from artefacts produced by the sensor or some part of the processing. Grey levels used should then be spread over enough of the range of the display device (computer screen or printer) to make this possible, without either making noise or artefacts too prominent or leaving some levels outside the range. If colour is used, the primary components should not be modified in a way that causes distracting changes in colour hues.

Grey levels, or colour components, are assumed to be non-negative integer values recorded in an array. Generally there will be an upper bound, the highest grey or component level produced by the sensor.

In this Report, sections 2, 3 and 4 each consider a different level of modification of a grey-level image (linear global, general global or local), discussing its basic method, its drawbacks and methods for avoiding those drawbacks. Section 5 discusses the extension of the method to colour, and Section 6 comments on non-linear effects of display devices and their impact on the material of earlier sections. Previous work is discussed under the relevant sections, and examples of applications are presented throughout. (Sections 2 and 3 describe existing techniques as a basis for comparison. The main contribution of the present work is in subsections 3.4 and 4.2).

2. Linear contrast enhancement

If a grey-level image is assumed to have its levels well distributed over a known range, it may be sufficient to apply a simple linear transformation to the level of each pixel. This approach is well known in image processing.

2.1 Scaling

In the simplest case of all, grey levels y in the range $[0, r]$ are mapped to display levels y' in the range $[0, G - 1]$, using

$$y' = \min(G - 1, \lfloor (G/r)y \rfloor).$$

If only a small part of the input range is used, only a correspondingly small part of the display range will be used, and contrast may be poor. Figure 1 shows an example of two parts of an infrared image with this problem, from the Goodrich DB-110 reconnaissance system.

2.2 Normalisation

If the input image has an actual grey level range $[p, q]$, found by searching for the minimum and maximum values, the appropriate mapping is

$$y' = \min(G-1, \lfloor G(y-p)/(q-p) \rfloor),$$

provided of course that $p \neq q$.

The main deficiency of this method is that it responds to the very lowest and highest grey levels without regard to how atypical they might be. In practice, the image may be affected by occasional noise spikes, or it may have some pixels set black or white along the edges because no data was available there. (Perhaps the sensor was a prototype.) There may also be white or black annotation on the image to give details of imaging location and conditions. In this case the grey levels used in the image may extend far beyond the range of meaningful data, and good contrast may not be achieved.

Figure 2 shows the result of normalising the DB-110 image portions. A few bad pixels have limited the contrast improvement.

2.3 Trimmed normalisation

The usual way to overcome the deficiency of normalisation is to ignore the bad values. As it is unlikely that they can be reliably distinguished automatically, the analyst must decide at least how many of the darkest and brightest pixels are unreliable. If the darkest $B_D n$ and the lightest $B_L n$ of n pixels are considered to be bad, the normalisation step must first find the quantiles $p = Q(B_D)$ and $q = Q(1 - B_L)$. (The quantile $Q(x)$ of a set of n values is defined to be a value that exceeds at most nx of the values and is exceeded by at most $n(1-x)$ of the values. As special cases, $Q(0)$ is the minimum, $Q(0.5)$ is the median and $Q(1)$ is the maximum.) The mapping is then

$$y' = \max\{0, \min(G-1, \lfloor G(y-p)/(q-p) \rfloor)\}$$

with allowance for the values outside $[p, q]$. For $p = q$ all values are mapped to an arbitrary constant such as $G/2$.

It is often safe to assume fixed values for B_D and B_L for a whole set of images, say 0.05 for each, to approximate the largest true values for the set. If the true values are sometimes much smaller, though, good pixels may be mapped to full black or white and useful details may be lost. Loss of contrast over most of an image will usually be a more serious error.

Figure 3 shows the result of trimmed normalisation of the DB-110 image portions, with $B_D = B_L = 0.001$. A much better range has been made visible in each portion. Details are poorest in the thin cloud in the first portion and in the extensive dark areas in the second.

3. Enhancement by histogram modification

3.1 General approach

The (normalised) histogram h of a discrete-valued image may be defined as a function of grey level with $h(y)$ being the fraction of pixels with grey level y , for $0 \leq y \leq G-1$. Equivalently, H is the cumulative histogram where $H(y)$ is the fraction of pixels with grey level below y , for $0 \leq y \leq G$. Then

$$\begin{aligned} H(y) &= \sum_{i < y} h(i) \\ H(0) &= 0 \\ H(G) &= 1 \end{aligned}$$

The shape of the histogram is sometimes considered as an indication of the quality of contrast in an image, with constant h (linear H) being the ideal. (Other functions for the ideal h including the normal probability function are also used but less often.) Clearly there are cases where the "ideal" is not appropriate - a clean image of a chessboard should have only two widely separated grey levels.

If a functional transformation $y' = f(y)$ is made to the grey levels, new histograms h' and H' result, with

$$\begin{aligned} H'(y') &= H(f^{-1}(y')) \\ f^{-1}(y') &= \min_{f(y) \geq y'} y \end{aligned}$$

Then contrast can be improved, by the quality criterion above, if $f(y)$ is chosen to make the histogram more like the ideal. (The discrete nature of grey levels will usually make an exact match impossible.)

3.2 Histogram equalisation

If constant $h'(i)$ is selected as the ideal histogram, it can be achieved approximately by the transformation

$$f(y) = \min(G-1, \lfloor GH(y) \rfloor) \quad (1)$$

for G grey levels in the output and any number in the input. Some distinct input grey levels may be mapped to the same output level, and some output levels may not be used, but repeating the transformation will make no further change.

Figure 4 shows the result of this transformation on the DB-110 image portions. The transformation often but not always improves local contrast.

Larson *et al* (1997), in their work on displaying images from sensors that record radiance that varies over several orders of magnitude, pointed out that histogram equalisation may perform poorly when a scene has a large area of constant radiance affected by small measurement noise. Because the area is large, its grey levels are stretched over a large part of the output range, and the noise is amplified to unacceptable levels. The author did not consider a related problem that occurs when a large sub-range of the input grey range is (almost) unoccupied: that sub-range is compressed in the output and areas with grey levels near either end of the sub-range may be poorly distinguished in the output. Clearly some adjustment is needed to ensure that the mapping $y' = f(y)$ is neither too steep nor too flat.

3.3 Bounded histogram equalisation

Larson *et al* (1997) offered a simple way to limit the steepness of the mapping $y' = f(y)$. They modified the histogram h by replacing values above a "ceiling" C by C itself, normalising the result to restore the total to 1, and iterating. In most cases, this produces a modified histogram h_b in which no value exceeds C , and values less than C are proportional to the original values in h . Other cases, in which too many or too few grey levels occur, must be handled separately.

Once a suitable h_b is available, the corresponding H_b can be calculated and used in place of H in (1). Over a short range of grey levels, the average gradient of $f(y)$, which represents the amplification of faint detail, is then no more than GC .

Now consider the case where the local average amplification is to be limited to the range $[A_F, A_C]$, where $A_F < 1 < A_C$. This requires a bounded histogram

$$A_F / G = F \leq h_b \leq C = A_C / G$$

with a "floor" as well as a ceiling. The iteration described above cannot be easily adapted to handle all cases, so a new approach is required.

3.4 Optimal bounded histogram equalisation

The problem of modifying the histogram to meet two bounds can be cast as a quadratic programming problem with coordinates $H_b(y)$. These should be as close as possible to the $H(y)$, within constraints on the first differences. A sum of square differences is suitable as the objective because it takes into account any overall shift between h and

h_B . (A sum of absolute differences could also be used. The problem is then a linear programming problem with separate positive and negative parts for each component difference. This objective relates to the road-builder's problem of how to change a given land profile to a given road profile with the same mean height, while shifting material from cuttings to embankments with the minimum amount of transport.)

The objective to be minimised is

$$P = \sum_{i=0}^G (H_B(i) - H(i))^2 \quad (2)$$

subject to

$$\begin{aligned} F &\leq H_B(i) - H_B(i-1) \leq C, \quad 1 \leq i \leq G \\ H_B(0) &= 0, \quad H_B(G) = 1. \end{aligned} \quad (3)$$

The objective is convex, and the solution will be as close as possible (in Euclidean distance) to the unconstrained solution H . The constraints are also convex, being the intersection of G infinite "slabs" and two hypersubplanes. The uniform distribution $H_U(i) = i/G$ is always feasible, because its first differences always meet the constraints for allowed A_F and A_G . This is a special case of quadratic programming and admits a relatively simple method of solution, at least approximately. An outline of the method follows.

Let H_C be the current best feasible estimate of H_B , initially set equal to H_U . Consider possible H_B along a straight line from H_C to H . Each first difference of H_B varies linearly with the fraction of the distance covered, so compute what fraction of the distance must be covered before each difference reaches one end of the allowed interval. (The result may be infinite; consider it greater than 1 in that case. Ignore the negative solution.) Find the minimum fraction over all differences. If this minimum is 1 or more, the unconstrained minimum H is feasible, $H_B = H$ and we are done.

If the minimum was less than 1, some constraint becomes active (i.e., satisfied with equality). Keep that constraint active; then some difference must be fixed, either at the floor or at the ceiling. Set H_C to the point where this happens. Assume that a constraint that becomes active remains active. (The assumption is not always correct, and will be considered later.) The approach to H will now be repeated, but with one difference fixed; under this constraint, the closest approach to matching H_C to H requires matching the mean of the two affected components in H_C to their mean in H . We may expect that further constraints will become active, and that we will be considering means of larger blocks of components of H_C and H , so we should represent these functions as blocks of values from the outset.

The initialisation and first step are now re-described in terms that can be used for later steps too. The $G+1$ coordinates are divided into contiguous blocks, initially one

coordinate in each block. H_C is set to H_U . Each block mean in H_C and H is computed. Consider H_B formed by varying the block means of H_C along a straight line towards the means of H , applying any change in a block mean to each coordinate in the block. A constraint can become active between adjacent blocks if the difference between the last coordinate of the first block and the first coordinate of the second block reaches the floor or ceiling. Find the smallest fraction of the change from H_C to H that makes this happen. Apply that fraction of the change (or the whole change if the fraction exceeds 1) to the block means of H_C , and change each coordinate by the mean change of its block. If the fraction exceeds 1, we are done. Otherwise make the new constraint active by merging two blocks and finding their combined means in both H_C and H . Then repeat the process until no more merges are possible.

An adjustment is needed to impose the end conditions in (3). In the first step, the conditions are automatic because the means for the first and last blocks are the same for H_C and H . If, in any step, an end block is merged with its neighbour, the mean of the new end block must be frozen so that the end coordinate does not change. This is achieved by setting the combined mean in H to match the combined mean in H_C . So long as coordinate values in H are not adjusted this has no side effects. When only two blocks remain, no further change is possible so the process terminates.

The above process terminates when the number of blocks is reduced to 2, or earlier if the block means of H_C and H can be made identical. The computation time for one step with n blocks is $O(n)$, so in the worst case with G grey levels the time is $O(G^2)$. While a single step may suffice, it was found in practice that most of the worst-case time is often required. If there are more than, say, 1024 grey levels, it may be necessary to reduce them to this number by normalisation first, to avoid excessive computation.

Now consider the possibility that the search has terminated but the solution can be improved by allowing an active constraint to become inactive again. This can be tested by splitting each block at each possible site in turn and testing whether an closer approach can be made to H with the blocks so modified. If an improvement is made, the split is retained and the approach to H is resumed. If a split allows no improvement it is undone and the next possible split is tried. If no split allows an improvement, then the solution is considered final. (In practice it appears that allowing constraints to be de-activated does not greatly alter the computation time or the resulting image. The possibility of improvement by de-activating more than one constraint at a time remains open. An exact but more elaborate method for solving the problem is given by Stoer (1971).)

Figure 5 shows the result of the process, with an amplification range 0.25 to 3, on the DB-110 image portions of Figure 1. The main effect of bounding the histogram in this case is a reduction of contrast. (This happens because the method avoids eliminating

the disused parts at the ends of the grey level range. It can be avoided by applying trimmed normalisation before the equalisation.) The Global Hawk image in Figure 10 better illustrates the effect; Figure 11 shows the results of equalisation after normalisation, with and without a ceiling.

4. Local contrast enhancement

If the contrast varies so much over an image that a simple mapping of grey levels is not enough, it may be better to modify the grey levels adaptively, adjusting them in a way that depends only on nearby values. This will mean that some coarser features of the image will be degraded as local details are enhanced (eg details in areas of shadow and bright sunlight will be equally visible, but it will be less obvious that the shadow was there). As long as a human analyst is aware of the possibility, and does not attempt to compare grey levels at widely separated points, the results will be useful. (Once interesting features have been detected, such comparisons should be done on a copy of the image enhanced by a different method.)

4.1 Possible approaches

If an image is divided into blocks and an enhancement method is applied to each block, local enhancement will be achieved, but block boundaries will probably be visible in the output. It will be better to use the moving-window approach, with each output value being computed from its own neighbourhood of pixels, and restricting the method to one that can be efficiently applied in this way.

Methods that can be efficiently applied in a moving window include trimmed normalisation (in which the output value depends on the input value and two quantiles in the window) and histogram equalisation (in which the output value depends on the rank of the input value within the window, determined by counting the smaller values). Much computing time can be saved for large window sizes in these methods by maintaining a moving window histogram and updating it appropriately when the window is moved by a single pixel.

Jobson *et al* (1997a) proposed a “retinex” method inspired by a theory of human visual processing. Essentially, a low-pass filter (in this case Gaussian) is applied to the image, then the logarithm of the ratio of the input to the filter output is calculated. A further linear transformation is needed for display. The scale of the filter is a matter of analyst preference, but the final transformation can be fixed for a wide variety of images.

Shortly afterwards, Jobson *et al* (1997b) sought to reduce the deficiencies of the retinex method with both small and large filter scales by applying the enhancement with three filter scales and taking the mean of the outputs. They achieved a balance between

showing coarse features correctly ("rendering") and emphasising fine features ("detail" and "dynamic range compression").

The moving window methods above did not do enough to preserve coarse features. The retinex method did this better, but fine features were not as well emphasised in bright areas as they might have been. A new method was sought to do both of these things better, even at the cost of fidelity to what human visual processing does, so that details are visible everywhere.

4.2 The ACE method

The following method is now proposed as a way to enhance fine detail in images of many grey levels and unpredictable contrast, while preserving enough of the coarse features to allow light and shade to assist in the recognition of objects:

1. A low-pass filter is applied to the image. The output is subtracted from the input, so that the image is expressed as the sum of low-pass (coarse) and high-pass (fine) components. (The fine component is allowed to take both positive and negative values.)
2. Trimmed normalisation is applied to the coarse component, so that inputs with many grey levels and unknown contrast can be handled nearly automatically, even in the presence of some extreme "bad" values, dark or bright. The result has the range $[0, G - 1]$, like a normalised full image.
3. The fine component f is equalised, i.e. given constant amplitude. This can be done by dividing f by the output \bar{f} of a low-pass filter applied to the absolute value of f . To avoid excessive amplification where only noise (or grey-level stepping) is present, the division is replaced by evaluation of $f' = \bar{f}f / (\bar{f}^2 + f_N^2)$, where f_N is a constant to be set approximately to the noise level. The result has a mean absolute value about 1, reduced where mainly noise is present. Use of the same filter as in step 1 gives satisfactory results.
4. The normalised coarse and equalised fine components are recombined with suitable weights.

The method has been dubbed ACE (for Adaptive Contrast Enhancement), and requires six parameters to control it. These parameters, with useful defaults, are:

1. The "bad pixel" fractions B_D and B_L for the coarse component. These will be increased by the low-pass filter. (Defaults both zero, on the assumption that pixels are good until suspected otherwise.)
2. The noise level f_N for the fine component. This will be hard to guess for a new imager with many grey levels. (Default 2, on the assumption that the grey level step is 1 and no noise is present, to suppress "contour map" patterns that result from emphasis of these steps.)

3. The weights for coarse and fine components in the final combination. (Defaults 0.75 for coarse, $G/4$ for fine.)
 4. The scale of the low-pass filter. This is method-dependent, and will be discussed below. (A filter resembling a Gaussian filter with $\sigma = 3$ is a satisfactory default.)
- (While six parameters may seem excessive for rapid viewing of a series of images, it is likely that values chosen for one image from a particular sensor can then be used for many images from that sensor without further adjustment.)

An efficient means of large-scale low-pass filtering is needed. Gaussian filtering allows separation into horizontal and vertical filters, which require computation time proportional to size rather than the square of size, but even more speed can be obtained by use of the following convolution kernels:

$$K_1 = K_2 = \begin{bmatrix} c & b & c \\ b & a & b \\ c & b & c \end{bmatrix}, K_3 = \begin{bmatrix} 0 & 0 & c & 0 & 0 \\ 0 & b & 0 & b & 0 \\ c & 0 & a & 0 & c \\ 0 & b & 0 & b & 0 \\ 0 & 0 & c & 0 & 0 \end{bmatrix}, K_4 = \begin{bmatrix} c & 0 & b & 0 & c \\ 0 & 0 & 0 & 0 & 0 \\ b & 0 & a & 0 & b \\ 0 & 0 & 0 & 0 & 0 \\ c & 0 & b & 0 & c \end{bmatrix}, \dots$$

Here $a = 1/4$, $b = 1/8$ and $c = 1/16$, so all kernels have nine taps with unit sum. After the second, each is produced by rotating the taps of the previous one by 45 degrees and increasing their spacing by $\sqrt{2}$. The x and y variances are $1/2, 1/2, 1, 2, 4, \dots$ but not all these kernels are low-pass. A low-pass filter of scale L is applied by convolving with K_1, K_2, \dots, K_L in turn, each being treated as sparse. The combinations have variances $1/2, 1, 2, 4, 8, \dots$ or standard deviations $2^{(L-1)/2}$, and the computation time is proportional to L , i.e. to $\log(\text{size})$. $L = 5$ is the recommended default scale.

ACE has been found useful on a wide range of imagery, including imagery with many grey levels that cannot be displayed directly. It is possible that the coarse component has too great a range for even coarse features to be displayed properly. In this case the weights for coarse and fine in the output can be set to, say, 0 and $G/2$. Some trace of the original light and shade pattern will remain near strong edges, and the analyst's visual system tends propagate these into other areas.

ACE allows the extension of Jobson *et al* (1997b), in that it can be applied with several filter scales (eg, $L = 5$, $L = 7$ and $L = 9$) and the average output taken. This is worth considering when haloes near major edges are distracting. At the same time, choosing a single filter scale to match the scale of blur in a sensor can produce a very desirable sharpening.

ACE can also be modified to cope with signal-dependent noise levels. In general it is assumed that the image is the form

$$(\text{coarse component}) + \sigma(\text{coarse component}) \times (\text{fine component}),$$

where $\sigma(\cdot)$ can be chosen to match the appropriate noise mechanism. It is then straightforward to extract the fine component and to recombine the modified components.

Figure 6 shows the DB-110 image portions of Figure 1 after processing by ACE. Figure 7 shows the effect of varying the noise level parameter. Figure 8 shows the effect of varying the weight for the coarse component. Figure 9 shows the effects of changing the filter scale and of averaging the outputs from three scales. Figure 10 shows the sharpening effects of the method on a Global Hawk EO image.

5. Application to colour imagery

5.1 Colour representations

There are many ways to represent colour in an image. The most common, RGB, specifies separate red, green and blue components, and corresponds to what is ultimately sent to a colour monitor. Other methods separate brightness from other colour characteristics, so that one colour parameter is sufficient to indicate whether a pixel would be light or dark in a monochrome version of the image.

While there are many arguments for or against using particular representations in colour image processing, one is very important when details are being enhanced. In the early days of colour monitors, human judgments of which RGB combinations were brighter or darker than others were tested. For the phosphors available at the time, the luminance $Y = 0.299R + 0.587G + 0.114B$ was considered to be the best normalised linear measure of the apparent brightness of a colour. The display is assumed to be balanced so that $R = G = B$ produces some shade of grey without noticeable colour tint; then a colour (R, G, B) has (Y, Y, Y) as the best grey approximation. The difference $(R - Y, G - Y, B - Y)$ is then the departure of the colour from grey, called the chrominance. It has also been noted that human vision is far less sensitive to fine details in the chrominance than to fine details in the luminance. (Limb *et al* (1977) have reviewed the history. The exact coefficients are given by Clarke (1991).)

It became standard, for television broadcasting purposes, to separate colour into luminance and chrominance, representing chrominance by two quantities which varied between TV standards (eg between NTSC and PAL). This allowed a monochrome receiver to receive a colour signal and display the luminance alone; it also allowed the chrominance to be transmitted with reduced resolution, to avoid exceeding the available bandwidths. Similar practices were taken up for videotape recording and in the JPEG compression technique for colour images, the same coefficients being used despite later changes in monitor phosphors. In some systems, the chrominance part of the signal seems to be more sensitive to noise and interference, too.

In view of the separate and unequal handling of luminance and chrominance in much colour imagery, it will often be unsafe to enhance red, green and blue components separately (because they all contain contributions from a poor-quality chrominance signal), or the chrominance itself. Rather, the luminance should be enhanced and the chrominance, in whatever form, left as it is. For this purpose, there is the possibility of using the ratio $(R/Y, G/Y, B/Y)$ instead of the three-component difference above.

5.2 Extension of grey-level methods

Jobson *et al* (1997b) recognised that their retinex method did not work well when applied to red, green and blue components of colour images separately. The problem was not with unreliable chrominance, but with different changes of contrast and the resulting colour fading. They corrected for it by adding a separate step that does, across the three colour bands, what the monochrome retinex enhancement does across pixels. This step deliberately exaggerates colour differences, but they found the effect acceptable. It even compensates for changes in the illumination source (e.g. from sunlight to blue sky or tungsten filament), but assumes that the scene is grey overall.

Any grey level enhancement method can be extended to colour images, including those with noisy or unreliable chrominance, as follows:

1. For each pixel compute Y from (R, G, B) according to the standard.
2. Treat the Y values as a monochrome image and find the enhanced values Y' .
3. For each pixel, compute either $(R + Y' - Y, G + Y' - Y, B + Y' - Y)$ or, if the alternative form of chrominance is to be used, $(RY'/Y, GY'/Y, BY'/Y)$.

Then the luminance is fully enhanced, while the chrominance (of the chosen form) is unchanged.

The choice of chrominance definition will affect the results where the luminance changes a lot. In particular, when dark areas become less dark, they will lose colour under the standard additive definition, but retain it under the alternative multiplicative definition. If poor colour settings have added a faint colour "cast" to dark areas, retaining the colour may give conspicuous bursts of unexpected hue; otherwise it can be an advantage. (Neither approach can do anything to relieve the atmospheric haze of a distant view, which fades the whole scene towards a light blue colour. That effect must be measured and removed by some separate mechanism first.)

Figures 12 and 13 show the application of histogram equalisation and ACE to colour images from a colour airborne reconnaissance system, using the multiplicative definition of chrominance.

Tests performed on JPEG images with strong compression have confirmed that enhancing only the luminance often gives greatly improved results.

6. Non-linear effects of display devices

In previous sections it was assumed that, when an image is displayed (by monitor, printer, on-line camera etc.), each pixel is viewed with a luminance proportional to its grey level. Real output devices do not always respond linearly to grey levels, and sometimes they affect the visibility of detail in the lightest or darkest areas. This report advocates manipulating contrast to improve visibility, and does not therefore consider accurate reproduction of colour. It is appropriate, however, to consider allowance for display non-linearity when it has an adverse effect.

The response z of a monochrome TV monitor to an input voltage V is often approximated by a "gamma law"

$$z = C \max(V + B, 0)^\gamma$$

where C depends on the contrast setting, B depends on the brightness setting and γ is a constant "gamma", ideally 1 for a linear response, typically about 2 for the tube itself, and possibly reduced for the monitor as a whole by compensating circuitry. In computer applications, the response to a grey level y out of G levels is then approximately

$$z = z_{\max} \{\max(y - y_0, 0) / (G - 1 - y_0)\}^\gamma$$

where y_0 is a threshold (usually negligible) and z_{\max} is the maximum response. Levels near or below the threshold are then displayed with much-reduced contrast.

For grey levels to be displayed linearly on a monitor, the non-linear effect can be approximately reversed by transforming each grey level y to an adjusted grey level y' , at or above the threshold, so that

$$\{(y' - y_0) / (G - 1 - y_0)\}^\gamma = y / (G - 1).$$

This requires

$$y' = y_0 + (G - 1 - y_0) \{y / (G - 1)\}^{1/\gamma}$$

which is described as "gamma correction" with the same gamma value. Image display software often includes an optional gamma correction of this kind. Whether a given monitor requires the correction, and how much, are matters of checking with test images.

For a colour monitor, each component may require its own correction. If the thresholds are negligible and the gamma values are similar it will be sufficient to apply the correction to the luminance, at least under the multiplicative definition of chrominance described in the previous section. For colour television there is an additional complication: transmission and recording standards require a fixed gamma correction for final viewing to be applied to the RGB signal before it is encoded (i.e. before luminance and chrominance are separated). Some attention must be given to how a TV image was extracted from a broadcast or a recording medium, in case an existing correction needs to be undone or taken into account.

Printing, especially in colour, is a complex process. For more modern digital printers, colour corrections may be performed internally, by their device drivers or by other colour management system (CMS) functions of operating systems. The outcome is complicated by adjustments in the raster image processor (RIP) in the device driver, which maps pixels to dots on the page. A check is needed in each case, and a gamma correction might still be useful.

On-line cameras typically involve a small TV monitor, colour filters and a photographic camera with film. The gamma law response of the monitor, the (approximately) logarithmic response of the film to the product of luminance and exposure time, and a possible further reversed logarithmic response if a negative film is printed, all contribute to the final output response. Experimentation is required. This method of output may not be very practical for surveillance work where speed is required.

Finally, the response of the observer's eye is important - changing the lighting level of a room can make a considerable difference to the apparent response of a monitor.

7. Conclusions

A range of techniques exists for improving the contrast of an image, both overall and for fine detail. These are useful when the contrast is poor, and when the image has too many grey levels or colours to be displayed directly. They range from simple scaling of values, through automatic adjustment of the range of values displayed, then general global mappings of brightness, to adaptive schemes that show local details over the whole image that would not be simultaneously visible under any global adjustment.

Two new methods, the constrained histogram equalisation method and the Adaptive Contrast Enhancement (ACE) method have been introduced for monochrome image enhancement. Two simple methods for extending monochrome methods to colour, while allowing for the deficiencies of some colour imagery, have also been described. Combinations of these methods will be useful for a wide range of surveillance imagery.

Further adjustments may be required to allow for non-linear responses of particular display devices, especially monitors.

References

Clarke C.K.P. (1991) Colour encoding and Decoding Systems, in Jackson K.G. & Townsend G.B. (eds), TV and Video Engineer's Reference Book, Butterworth-Heinemann Ltd., Oxford, 11/1-11/22.

Jobson D.J., Rahman Z. & Woodell G.A. (1997a) Properties and Performance of a Center/Surround Retinex, *IEEE Transactions on Image Processing*, Vol. 6 No. 3, 451-462.

Jobson D.J., Rahman Z. & Woodell G.A. (1997b) A Multiscale Retinex for Bridging the Gap Between Color Images and the Human Observation of Scenes, *IEEE Transactions on Image Processing*, Vol. 6 No. 7, 965-976.

Larson G.W., Rushmeier H. & Piatko C. (1997) A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes, *IEEE Transactions on Visualization and Computer Graphics*, Vol. 3 No. 4, 291-306.

Limb J.O., Rubinstein C.B. and Thompson J.E. (1977) Digital coding of Color Video Signals - A Review, *IEEE Transactions on Communication*, Vol. 25, 1349-1384.

Stoer J. (1971) On the Numerical Solution of Constrained Least-squares Problems, *SIAM Journal of Numerical Analysis*, Vol. 8 No. 2, 382-411.

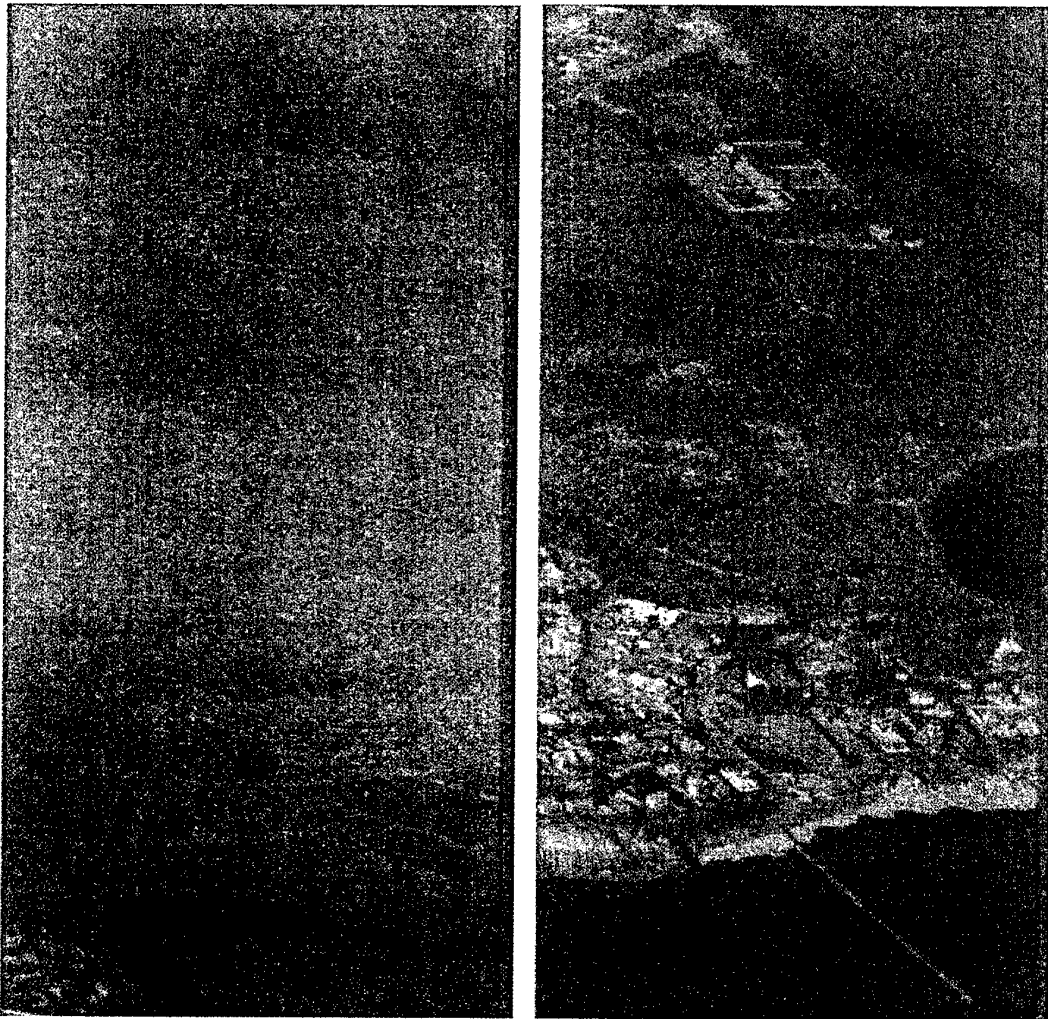


Figure 1. Two portions of an infrared image of Port Noarlunga, South Australia, from the Goodrich DB-110 reconnaissance system. The original was a 512x5428 12-bit image, formed as a mosaic of sequential frames whose joins are visible in enhanced versions. The standard deviations of the grey levels in these portions were only 65 and 122. Here the grey levels have been divided by 16 for 8-bit viewing, and most of the grey level range is little used.

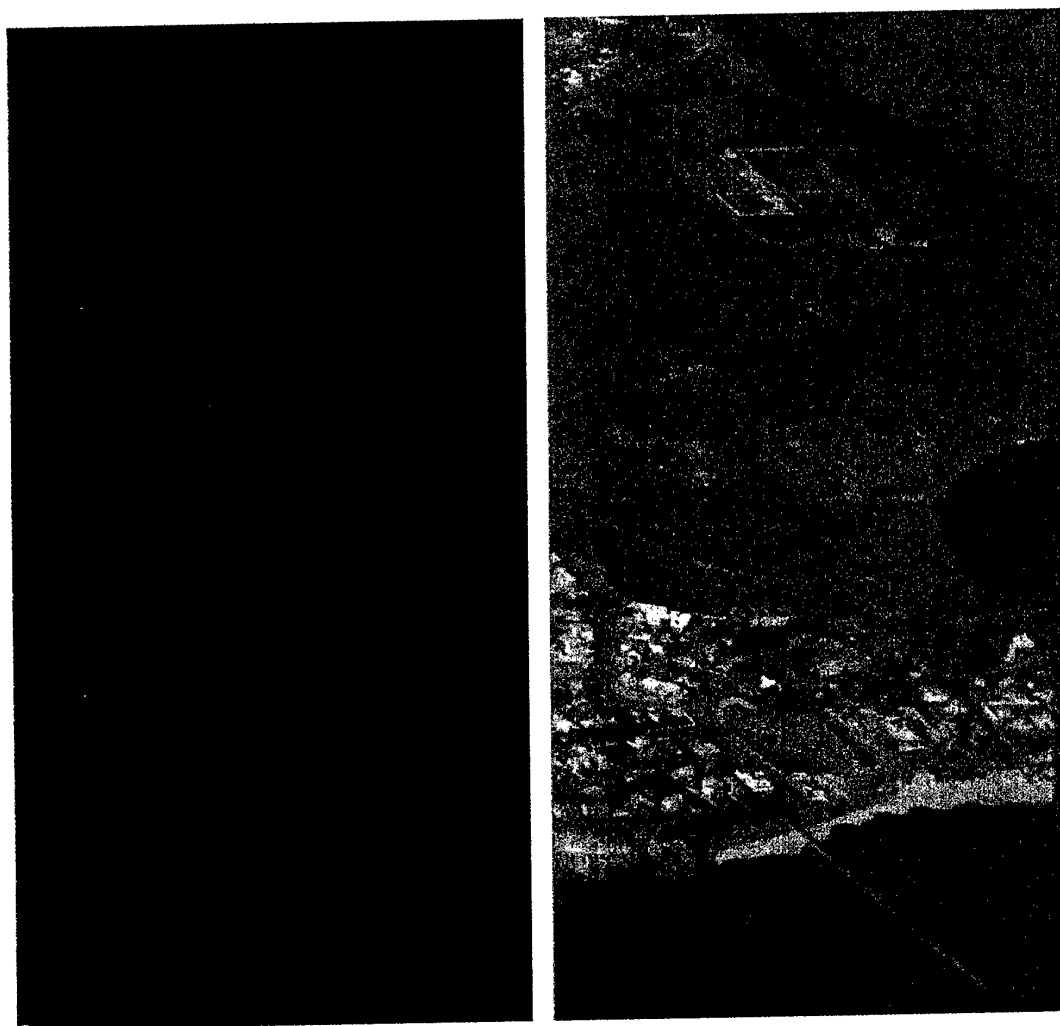


Figure 2. The DB-110 image portions converted separately from 12 to 8 bits using normalisation. Stray large and small values have misled the method and there is little improvement over division by 16.

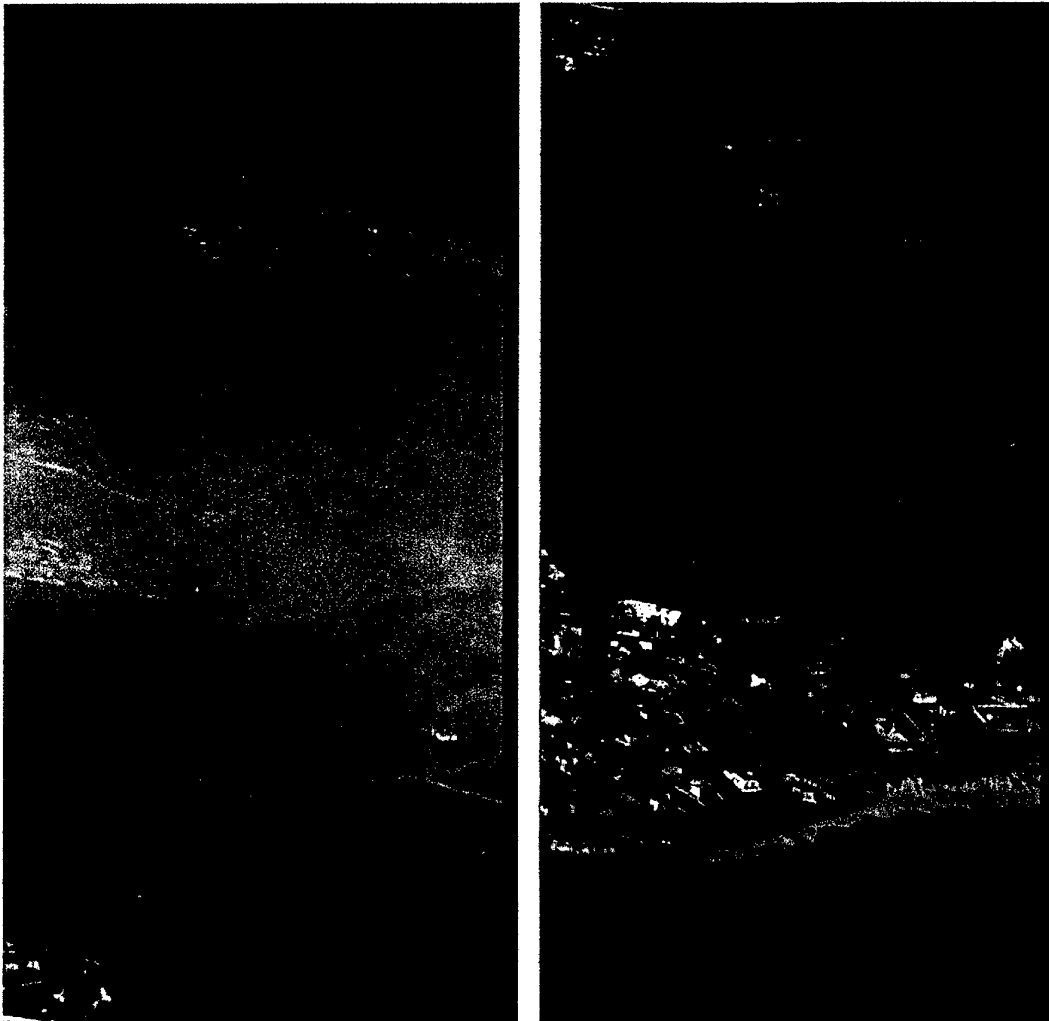


Figure 3. The DB-110 image portions converted separately from 12 to 8 bits by trimmed normalisation, treating the darkest 0.1% and brightest 0.1% of pixels as bad. The extended light areas in the first portion are thin cloud. Much of the second portion is too dark to show details.

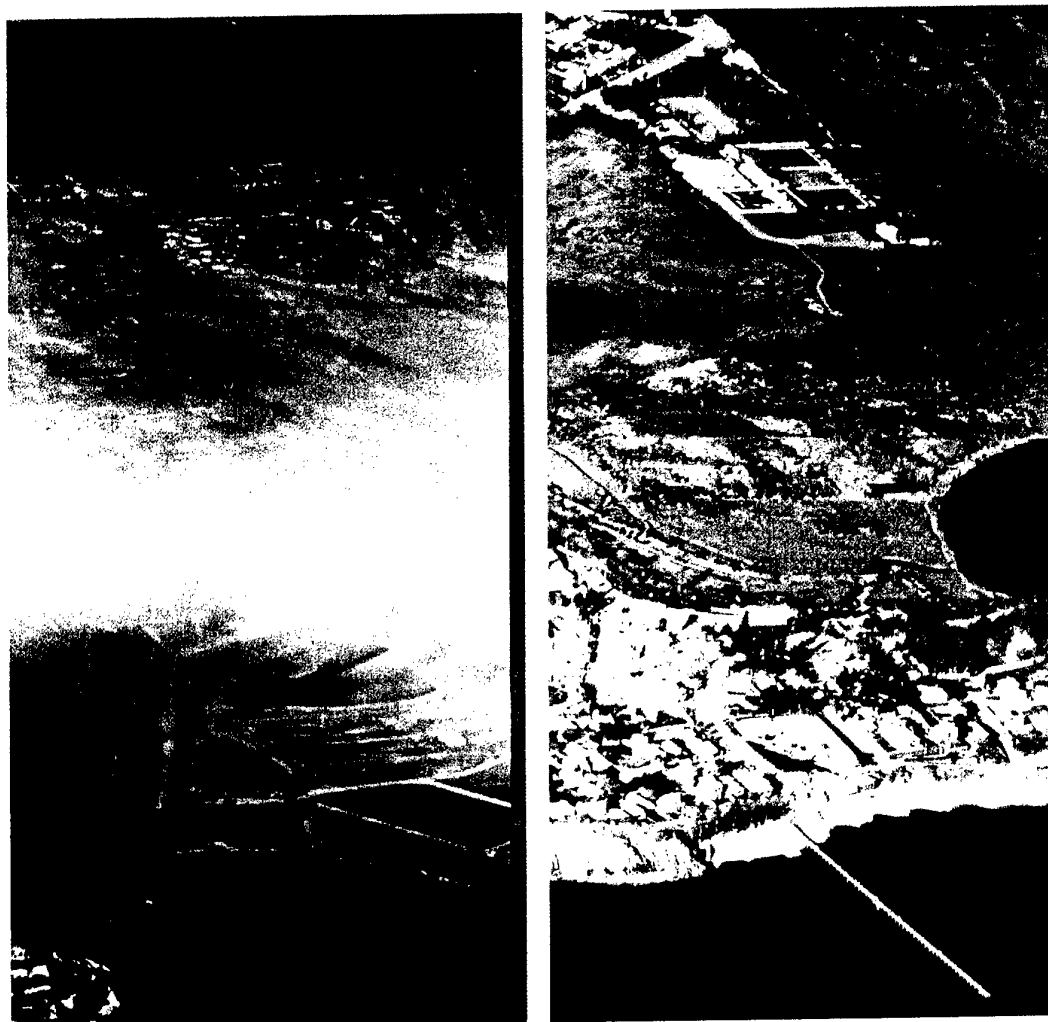


Figure 4. The DB-110 image portions after separate histogram equalisation (with 4096 grey levels) then division by 16. The second portion has much more useful contrast, but the first still has extensive areas of low contrast.

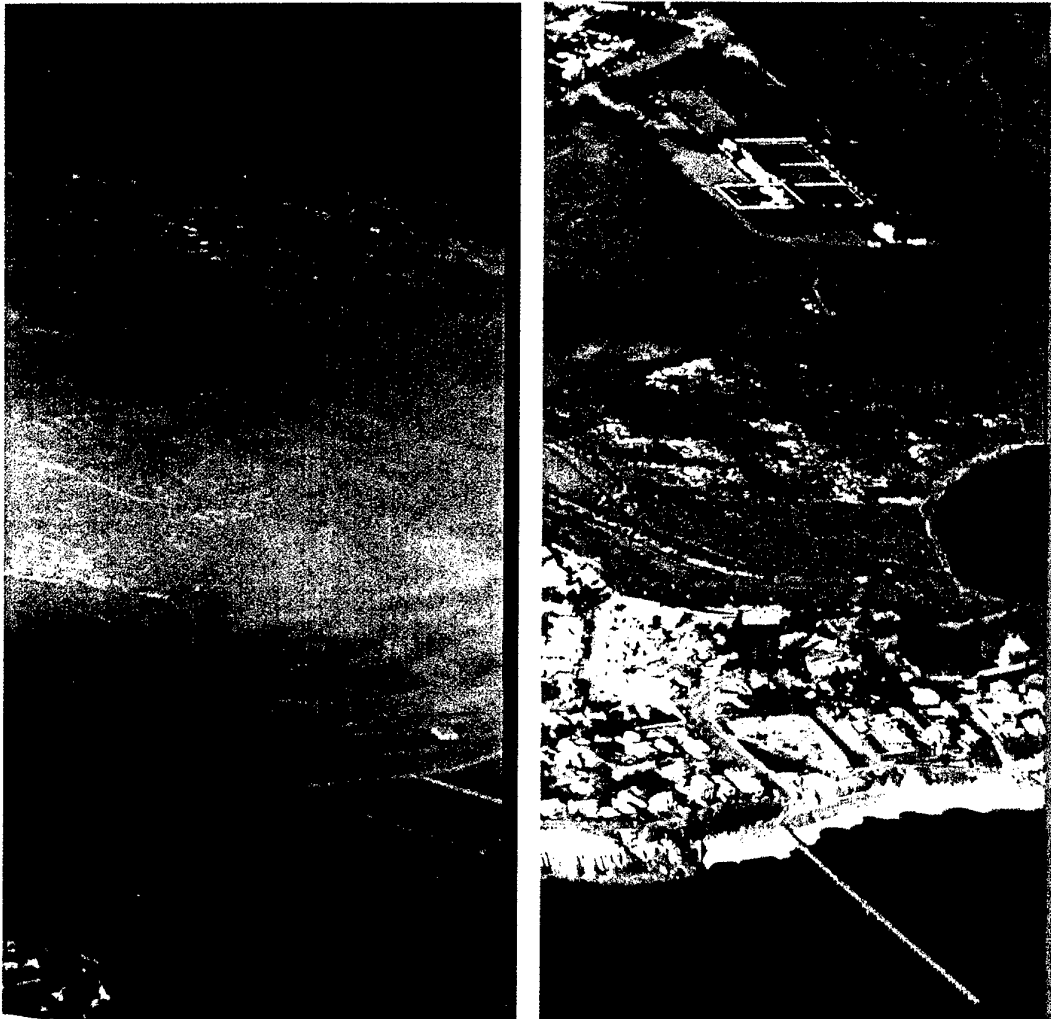


Figure 5. The DB-110 image portions after separate bounded histogram equalisation with a floor of 0.25 and a ceiling of 3. The limited grey level range has reduced the contrast in this case, relative to the results of full unbounded histogram equalisation.

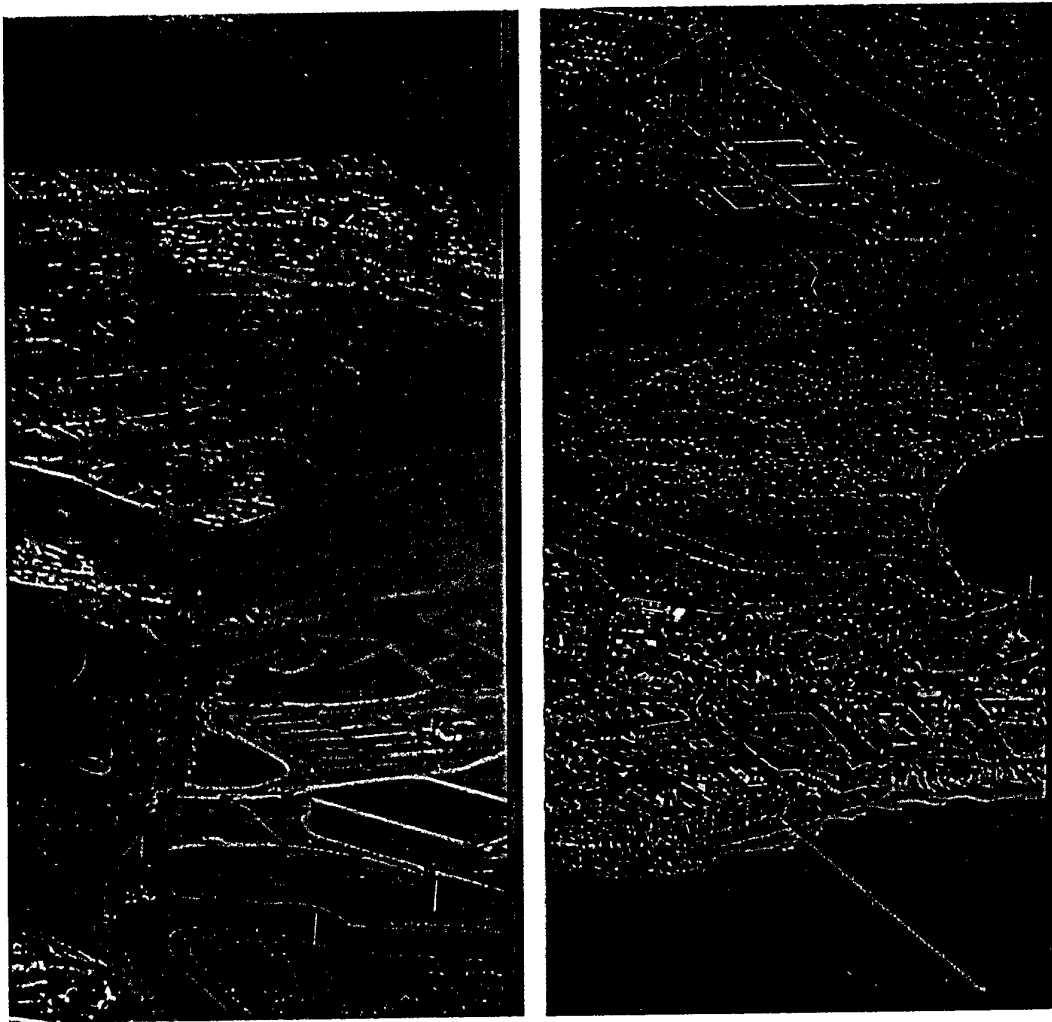


Figure 6. The DB-110 image portions after processing by ACE, with a suitable noise level. Details with higher contrast than the noise level have been revealed, including sea waves and some detail under the thin cloud. Seams in the mosaic are now more evident. Grey levels cannot now be compared between areas – the surface of the jetty does not really change to darker material over the beach.



Figure 7. The effect of noise level on ACE processing. The top portion was produced with a level of 0.1 (i.e., assuming negligible noise) and reveals a CCD pixel calibration pattern that is normally not noticeable. The bottom portion was processed with a level of 5 and the pattern remains faint.

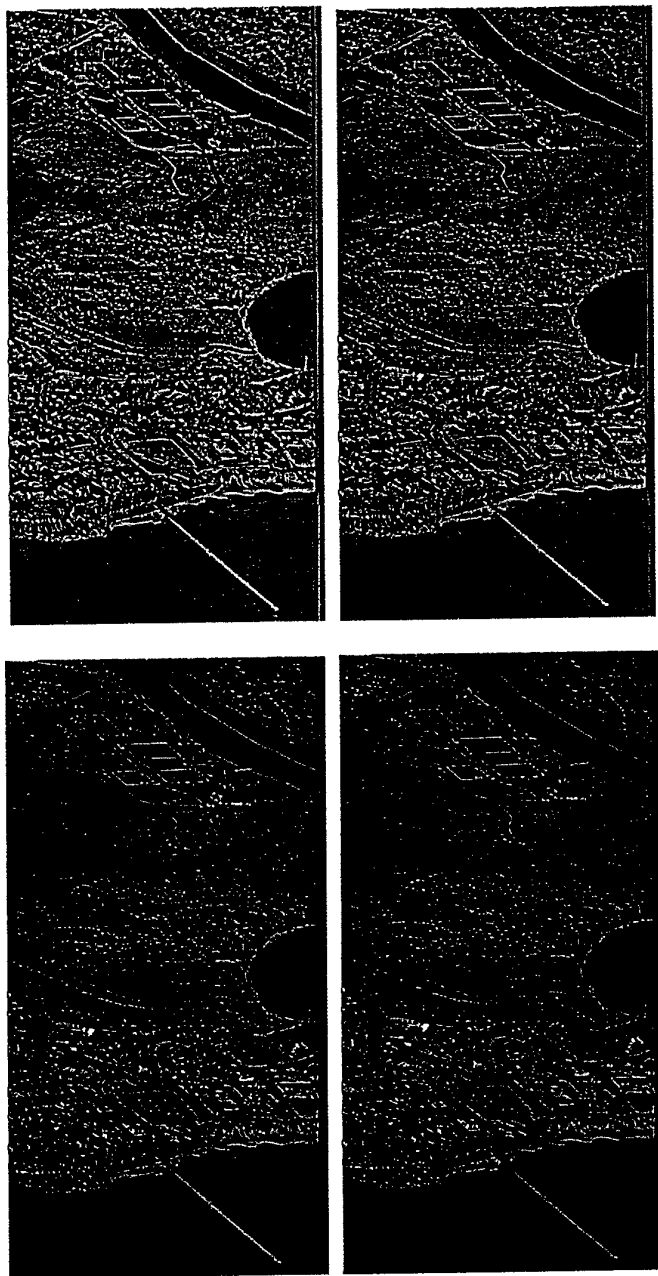


Figure 8. The effect of varying the weight of the normalised coarse component in ACE processing. The values used were 0, 0.25, 0.5 and 0.75. The coarse component aids object recognition but an image is often still intelligible without it.

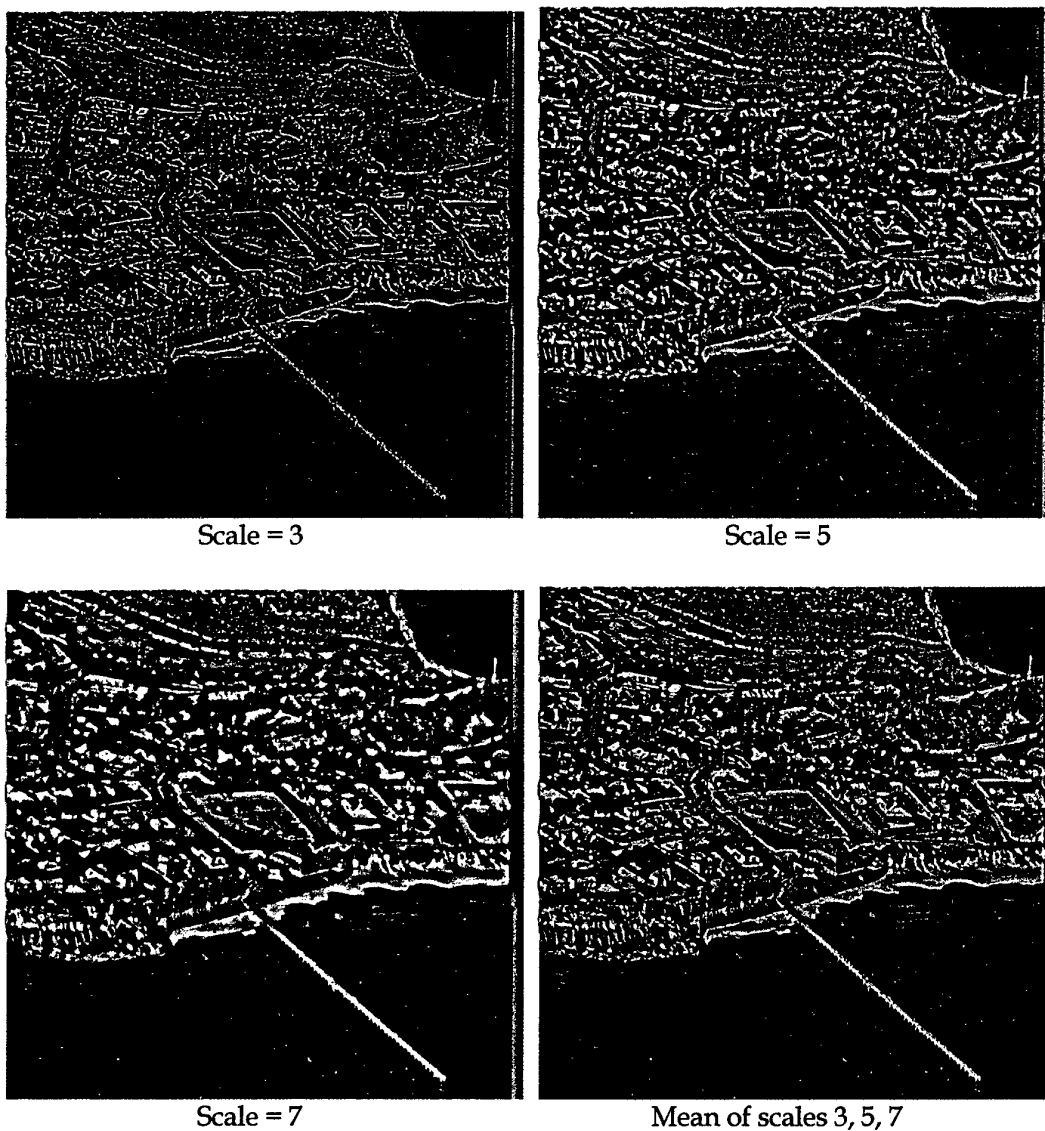
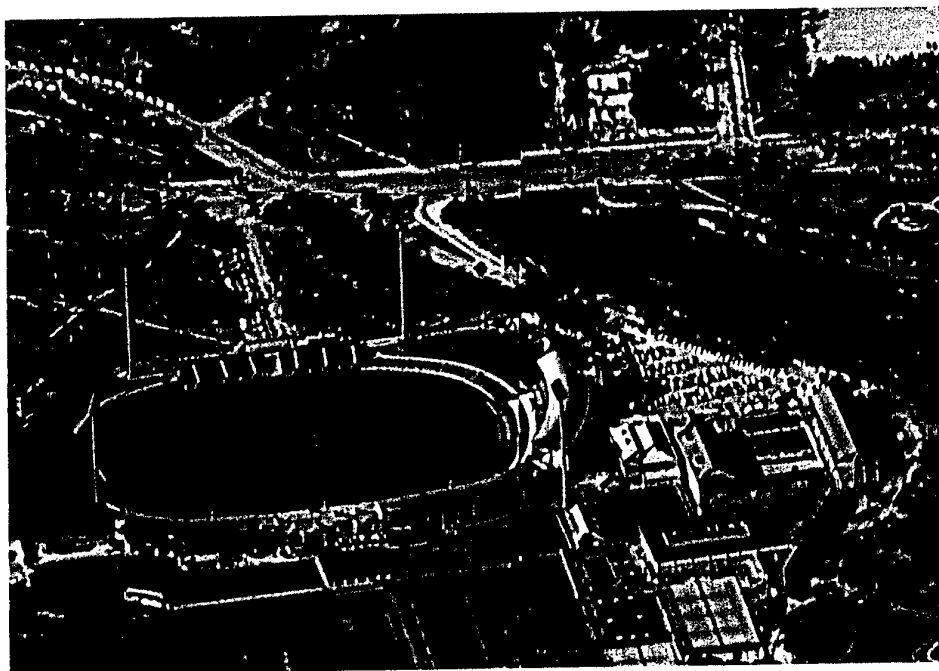


Figure 9. The effect of varying the scale in ACE processing. The mean of the outputs for three scales is not the same as the output for an intermediate scale and may reduce the distraction of haloes near strong edges. (Each scale step is a size increase of $\sqrt{2}$.)



Trimmed normalisation



ACE

Figure 10. Part of a Global Hawk EO mosaic image showing the Adelaide Oval, processed by trimmed normalisation alone (with 0.1% each of bad dark and bright pixels) and by ACE. This image is sharpened by the halo effect of ACE.

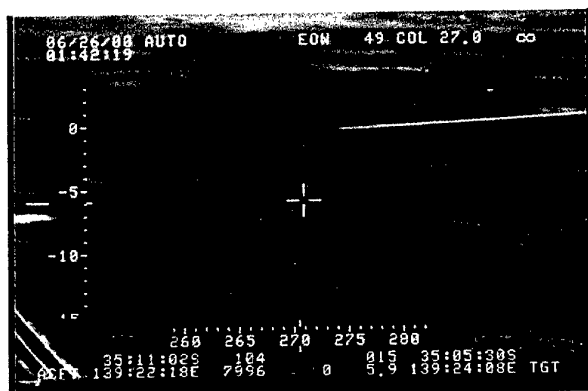


No bound

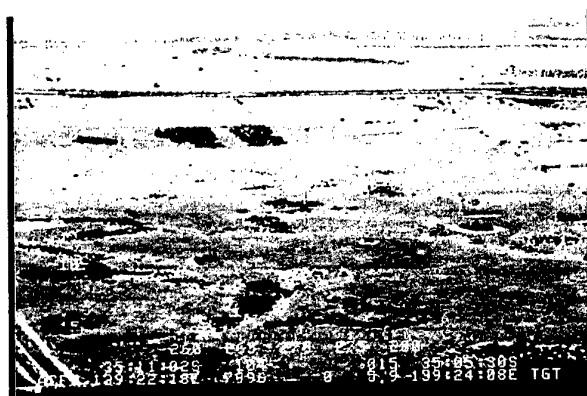


Ceiling = 3

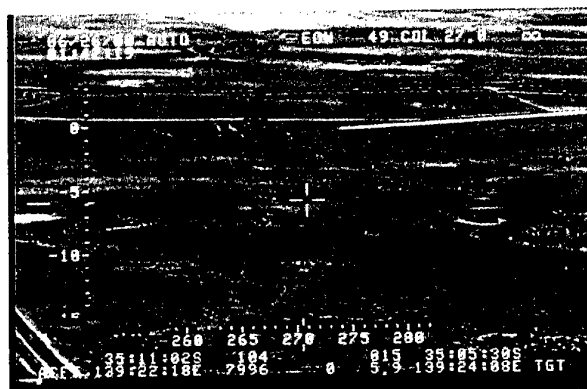
Figure 11. The Global Hawk image processed by normalisation followed by histogram equalisation, with and without bounding by a ceiling only. The effect of grey level steps on the playing area is exaggerated without the ceiling of 3.



Input image

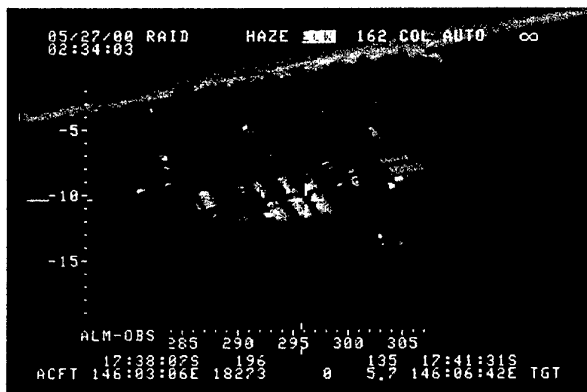


Bounded histogram equalisation



ACE

Figure 12. Application of bounded histogram equalisation and ACE to a colour airborne reconnaissance image from the Murray Bridge area, South Australia (adjusted by hand to remove haze). Only ACE was able to suppress apparent JPEG artefacts in this image



Input image



Bounded histogram equalisation



ACE

Figure 13. Application of bounded histogram equalisation and ACE to another colour airborne reconnaissance image from the Darwin area

DISTRIBUTION LIST

Enhancing Imagery to Improve the Visibility of Detail

Robert Whatmough

AUSTRALIA

DEFENCE ORGANISATION

	No. of copies
S&T Program	
Chief Defence Scientist	shared copy
FAS Science Policy	
AS Science Corporate Management	
Director General Science Policy Development	
Counsellor Defence Science, London	Doc Data Sheet
Counsellor Defence Science, Washington	Doc Data Sheet
Scientific Adviser to MRDC, Thailand	Doc Data Sheet
Scientific Adviser Joint	1
Navy Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser - Army	Doc Data Sht & Dist List
Air Force Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser to the DMO M&A	Doc Data Sht & Dist List
Scientific Adviser to the DMO ELL	Doc Data Sht & Dist List
Director of Trials	Doc Data Sht & Dist List
Information Sciences Laboratory	
Chief of Intelligence Surveillance and Reconnaissance Division	Doc Data Sht & Dist List
Research Leader Imagery Systems	1
Head Image Analysis and Exploitation Group	1
Mr. Don Uksi	1
Dr. Tim Miller	1
Dr. Nick Redding	1
Ms. Merryll Fiebig	1
Mr. Geoffrey Burls	1
Mr. Peter Virgo	1
Mr. Robert Whatmough	6
DSTO Library and Archives	
Library Edinburgh 1 & Doc Data Sheet	
Australian Archives	1
Capability Systems Division	
Director General Maritime Development	Doc Data Sheet
Director General Aerospace Development	Doc Data Sheet
Director General Information Capability Development	Doc Data Sheet
Office of the Chief Information Officer	

Deputy CIO	Doc Data Sheet
Director General Information Policy and Plans	Doc Data Sheet
AS Information Structures and Futures	Doc Data Sheet
AS Information Architecture and Management	Doc Data Sheet
Director General Australian Defence Simulation Office	Doc Data Sheet

Strategy Group

Director General Military Strategy	Doc Data Sheet
Director General Preparedness	Doc Data Sheet

HQAST

SO (Science) (ASJIC)	Doc Data Sheet
----------------------	----------------

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW	Doc Data Sht & Dist List
Director General Navy Capability, Performance and Plans, Navy Headquarters	
Doc Data Sheet	
Director General Navy Strategic Policy and Futures, Navy Headquarters	
Doc Data Sheet	

Air Force

SO (Science) - Headquarters Air Combat Group, RAAF Base, Williamtown	
NSW 2314	Doc Data Sht & Exec
Summ	

Army

ABCA National Standardisation Officer, Land Warfare Development Sector,	
Puckapunyal	e-mailed Doc Data Sheet
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD	
Doc Data Sheet	
SO (Science) - Land Headquarters (LHQ), Victoria Barracks NSW	
	Doc Data & Exec Summ

Intelligence Program

DGSTA Defence Intelligence Organisation	1
Manager, Information Centre, Defence Intelligence	
Organisation	1 (PDF version)
Assistant Secretary Corporate, Defence Imagery and Geospatial Organisation	
Doc Data Sheet	

Defence Materiel Organisation

Head Airborne Surveillance and Control	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Electronic Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Head Land Systems Division	Doc Data Sheet
Head Industry Division	Doc Data Sheet
Chief Joint Logistics Command	Doc Data Sheet
Management Information Systems Division	Doc Data Sheet
Head Materiel Finance	Doc Data Sheet

Defence Libraries

Library Manager, DLS-Canberra	Doc Data Sheet
Library Manager, DLS - Sydney West	Doc Data Sheet

OTHER ORGANISATIONS

National Library of Australia	1
NASA (Canberra)	1

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy	
Library	1
Head of Aerospace and Mechanical Engineering	1
Serials Section (M list), Deakin University Library, Geelong, VIC	1
Hargrave Library, Monash University	Doc Data Sheet
Librarian, Flinders University	1

OUTSIDE AUSTRALIA**INTERNATIONAL DEFENCE INFORMATION CENTRES**

US Defense Technical Information Center	2
UK Defence Research Information Centre	2
Canada Defence Scientific Information Service	e-mail link to pdf
NZ Defence Information Centre	1

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1

SPARES	5
--------	---

Total number of copies:	38 Paper & 1 pdf	39
--------------------------------	-----------------------------	-----------

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)			
2. TITLE Enhancing Imagery to Improve the Visibility of Detail				3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) <div style="text-align: right;"> Document (U) Title (U) Abstract (U) </div>			
4. AUTHOR(S) Robert Whatmough				5. CORPORATE AUTHOR Information Sciences Laboratory PO Box 1500 Edinburgh South Australia 5111 Australia			
6a. DSTO NUMBER DSTO-TR-1557		6b. AR NUMBER AR-013-039		6c. TYPE OF REPORT Technical Report		7. DOCUMENT DATE February 2004	
8. FILE NUMBER 9505/25/267		9. TASK NUMBER INT 02/235		10. TASK SPONSOR		11. NO. OF PAGES 29	
						12. NO. OF REFERENCES 6	
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1557.pdf				14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division			
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <div style="text-align: center;"><i>Approved for public release</i></div>							
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111							
16. DELIBERATE ANNOUNCEMENT No Limitations							
17. CITATION IN OTHER DOCUMENTS				Yes			
18. DEFTEST DESCRIPTORS Image enhancement, Image contrast, Gray scale, Surveillance imagery							
19. ABSTRACT Imagery from surveillance sensors is often not suitable for immediate viewing, because the number of grey levels or colours used is much larger or smaller than the number available on the display device. This Report considers some old and new techniques for making imagery ready to view, so that both broad features and fine detail are visible. The basic techniques often give poor results when they encounter the peculiarities of particular sensors or scenes; in these cases variations are available that overcome the problems at a cost in computing time. Examples of the application of techniques and variations to airborne sensor imagery are given.							